

# 3D Reconstruction, Segmentation and Classification of Corals from Aerial Images

## Final Report: CS 231A

Hardik Kabaria<sup>a</sup>, Ashish Goel<sup>b</sup>, Kristen Moore<sup>c</sup>

<sup>a</sup>*Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA, email:  
hardikk@stanford.edu*

<sup>b</sup>*Department of Aeronautics & Astronautics, Stanford University, Stanford, CA 94305, USA, email:  
ashish09@stanford.edu*

<sup>c</sup>*Department of Mathematics, Stanford University, Stanford, CA 94305, USA, email:  
klmoore@stanford.edu*

---

### Abstract

Coral reefs cover an area of over  $280,000\text{km}^2$  and support thousands of species in what many describe as the “rain forests of the seas”. Coral reefs face numerous threats, with current estimates suggesting that thirty percent of corals are in critical condition and may die within 10 to 20 years [cite]. One resourceful way to monitor the health of a coral reef is to periodically take aerial images of the reef from a UAV, use these images to construct a 3D model of the reef and monitor its evolution in time. Here we present a methodology for the 3D reconstruction, segmentation and classification of corals from aerial images taken by an UAV. The proposed methodology consists of building 3D point cloud from the images based on “structure from motion” approach, and then building the 3D surface to fit the point cloud. We segment the 3D surface based on the discrete mean curvature of the surface, cluster the coral containing segment into groups of individual corals using mean-shift algorithm and finally classify the corals into two classes of coral species namely, Branching and Porites using SVM. With the help of experts in the field of marine biology, we verify that clusters generated by our algorithm are indeed clusters of corals with high precision and recall. Using their labeling of our dataset, we achieve 76.6% success in the classification process. When compared to standard 2D techniques, we find that both the segmentation and classification processes benefit from using the 3D information.

---

### 1. Introduction

By increasing sea level and reducing coral growth, climate change is poised to have a huge impact on reefs and on human uses of coastlines. When coral growth stops, the reef framework decays, and waves reach the shore more often. This creates beach erosion, coastal retraction and destruction of built structures. The economic and social loss from such damage adds 100s of billions of dollars to other impacts of reef degradation such as biodiversity or fisheries loss. Management methods to slow the loss of reefs have stalled because of a fundamental but simple problem - accurate 3-D maps of coral reefs,

*Preprint submitted to CS 231A teaching staff*

*March 18, 2014*

distinguishing live from dead corals and high resolution analysis of coral growth potential, are largely lacking.

In order to address this, Stanford students Ved Chirayath, Tamaki Bieri and Trent Lukaczyk spent a few weeks during the 2013 summer at Ofu island in the American Samoa, developing and testing a novel method for aerial imaging of corals from unmanned aerial vehicles. By imaging aurally, they were able to scan through the entire 2km length of the coral reef, something which would not be possible by traditional diving based coral imaging techniques. Having collected these images, the fluid lensing technique developed by Ved Chirayath was applied to these images to eliminate the effects of wave distortions on the surface of water.

The objective of our team for the computer vision project was to take these images and use them to carry out 3D reconstruction of the reef, segregate the corals from non-coral objects in the images such as rocks, and finally to estimate the total volume of the corals in the reef, which is of immense value in the field of marine biology for monitoring the health of the corals. Since the aim of this project is to detail the approach to be taken in each step of the process in order to ultimately estimate the volume coral in a particular imaged region, we have used a relatively small subset of the images, rather than trying to implement our coral mapping program for the entire coral reef.



Figure 1: Image of the UAV taking aerial images of the coral reef in American Samoa.

The organization of the paper is as follows. First we present our methodology and results for generating the 3D point cloud and fitting to it a three dimensional surface. We present the verification results obtained after showing out output to the experts in the field and finally results from the classification of these corals into two different species, based on their labeling.

## 2. Methodology

### 2.1. Overview

Here we describe the different steps of our approach. We used external software to generate 3D point cloud from the images. We processed that point cloud with two different approaches in order to segment point cloud into regions with and without the corals. Next step, we divide the coral containing regions into individual corals. And as a final step we apply SVM based classification algorithm to classify the corals into two different species.

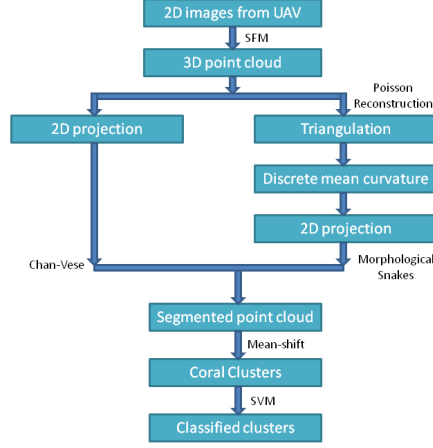


Figure 2: Overview of the methodology adopted.

### 2.2. Structure from motion

The 3D point cloud was generated using a commercial software package called PhotoScan by Agisoft, which implements the structure from motion technique. The point cloud is then used as the starting point of our algorithm for segmentation, clustering and classification.

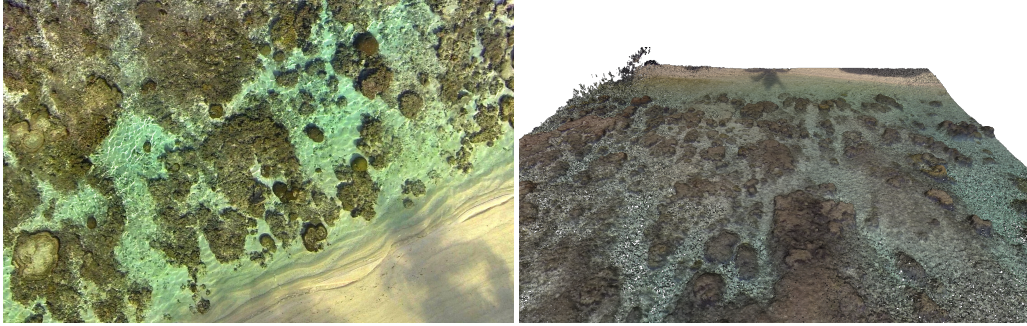


Figure 3: A sample of an aerial image taken and the 3D point cloud generated from such images based on structure from motion.

### 2.3. Segmentation

We implemented two different approaches to segment the point cloud into regions that contain corals and regions that don't. Both approaches are based on an algorithms designed for segmenting objects that don't have sharp edges. The idea is to start with a closed curve around the object to be detected in the given image, and move the curve in the direction of its inward pointing normal vector with a speed corresponding to its mean curvature plus a forcing term that's chosen to make the evolution stop on the boundary of the object.

### 2.3.1. Approach 1

Our first approach segmented the 3D point cloud based on the intensity of the pixels. Since the point cloud consists of 2.8 million points, we first projected the 3D point cloud on a 2D gray scale image of  $512 \times 512$  pixels, and then applied an algorithm developed by Chan and Vese in [1]. In particular, the zero level-set of the function  $\phi(x, t)$  describing the curve at time  $t$  is evolved according to

$$\frac{\partial \phi}{\partial t} = \text{div} \left( \frac{\nabla \phi}{|\nabla \phi|} \right) - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2, \quad (1)$$

where the first term on the right hand side corresponds to the mean curvature of the zero level-set of  $\phi(x, t)$ , and the remaining stopping terms depend on the average  $c_1$  and  $c_2$  of the intensity  $u_0(x)$  of the image regions inside and outside the curve respectively. Chan and Vese named this the “active contours without edges” approach, since the stopping term isn’t an edge function depending on the image gradient, as is the case in geometric active contour approaches.

An example of a segmented image is shown in Figure 4, and the video of the segmentation evolution can be seen here <https://www.dropbox.com/s/u0b96gw0eyu00ks/CVmovie.wmv>.

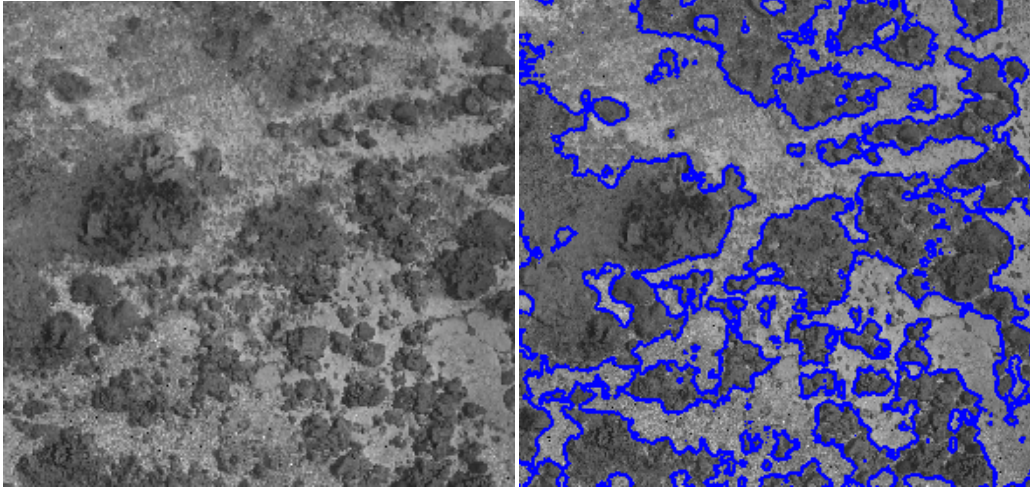


Figure 4: Segmentation performed to segment coral containing regions from non-coral related pixels after converting the point cloud to a gray scale image.

As one can see from results in Fig. 4, the Chan-Vese algorithm does a pretty good job of segmenting the regions containing corals and produces very few artifacts. However, a closer inspection would reveal that when the sea floor has a dark color, the algorithm still identifies it as a coral containing region. This motivates us to exploit the third dimension to achieve better segmentation results.

### 2.3.2. Approach 2

The 3D point cloud puts the geometry of the reef at our disposal, making the curvature of the 3D reef reconstruction a natural criteria for segmentation. To carry out a curvature-based segmentation, we first fit a 3D triangular surface consisting of only

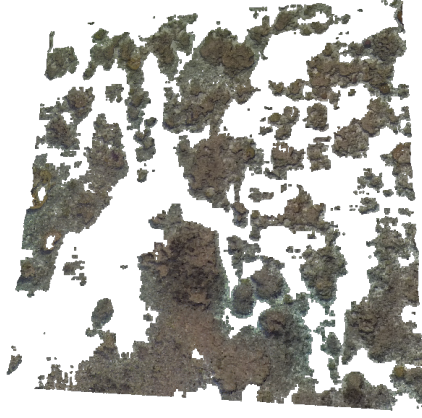


Figure 5: Output generated at the end of approach 1 which consists of 3D points corresponding to corals with original pixel colors.

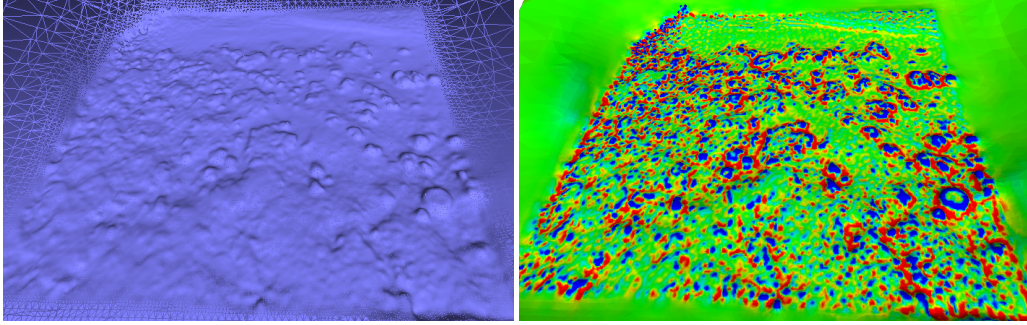


Figure 6: 3D triangular surface fitting based on Poisson reconstruction of the point cloud. The triangulation consists of 274067 triangles. The colors of the pixels on the right represent the mean curvature values.

274067 faces to the point cloud of 2.8 million points. This was done using the Poisson reconstruction approach of [2], which expresses surface reconstruction problem as the solution to a Poisson equation. The idea of this approach is to compute the indicator function  $\chi$  for the surface and then use that to extract an appropriate isosurface. This approach uses the observation that the gradient of the surface indicator function is a vector field  $\mathbf{v}$  that is zero almost everywhere, except at points near the surface, where it corresponds to the inner normal of the surface. The surface reconstruction problem then amounts to computing the scalar function  $\chi$  whose gradient best approximates the vector field  $\mathbf{v}$  defined by the point samples. This is equivalent to solving the Poisson problem

$$\Delta\chi = \text{div}(\mathbf{v}). \quad (2)$$

We used the computational geometry library CGAL [3] to perform the Poisson surface reconstruction, with the outcome shown in Figure 6.



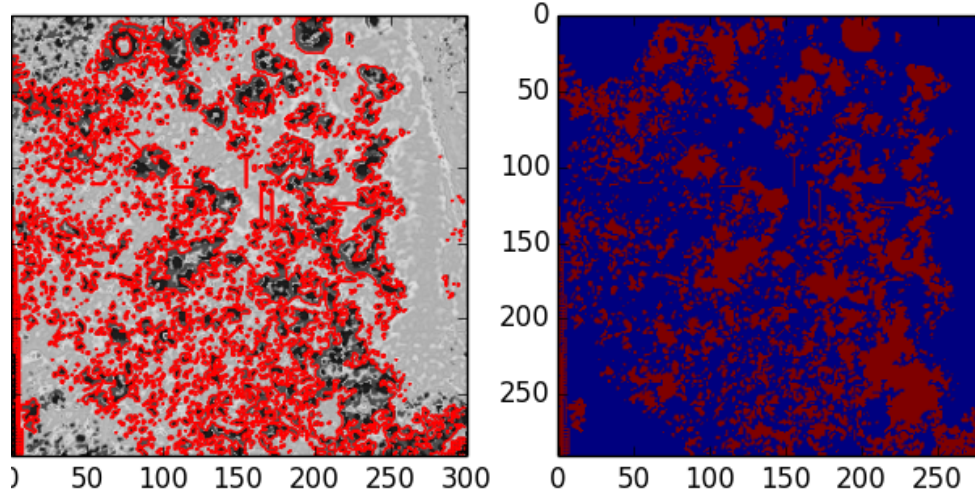


Figure 7: Binary segmentation performed on the gray scale image derived from the curvature colors.

The discrete mean curvature of this triangulated surface was then calculated using Meshlab, as depicted graphically with a color mapping on the triangulated surface in Figure 6. Preliminary inspection of Figure 6 suggests that the boundaries of the corals correspond to regions of high curvature, prompting the use of the curvature map to segment the corals from the sea bed.

In order to carry out this curvature based segmentation, we first projected the 3D curvature color map onto a 2D image and converted it to a gray scale image as shown in Figure 7. We then used the “morphological active contours without edges” algorithm, proposed by Pablo Mrquez-Neila et al in 2013 [4], to segment the 2D image. This morphological algorithm has equivalent infinitesimal behavior to (1), but instead of a PDE for the level set function, the morphological operator uses the sup and inf of the level-set function in very small balls around each point on the curve in order to evolve the curve towards the desired object. The video of the morphological segmentation can be seen here <https://www.dropbox.com/s/4a2se1uyfqk14d1/out.mp4>.

The output of the segmentation process indicates which points of the Poisson reconstructed surface in Figure 6 should be discarded. We can see that this method produces sharper boundaries and does a good job of removing the sea floor but produces several noisy regions. In order to fix this to some extent, we applied a smoothing filter to the output of the second approach. It reduces the sharpness of the cluster boundaries but also gets rid of the spurious, miniscule clusters. Simple thresholding by cluster size and merging small clusters with large neighboring clusters was also tried and the results were comparable.

We can also see that some of these corals have flat top surfaces which makes their

curvature resemble that of the sea floor producing holes in the coral-containing segments. This can be fixed by using not just the curvature but the height of the points in the original 3D point cloud in the segmentation process.

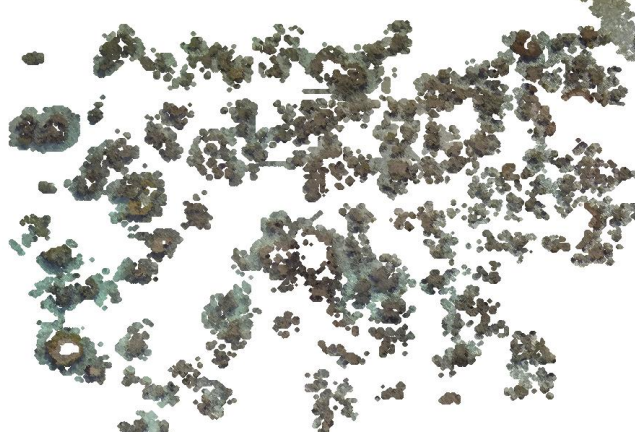


Figure 8: Output of approach 2 consisting of 3D points with original pixel colors in regions corresponding to corals.

#### 2.4. Clustering and Verification

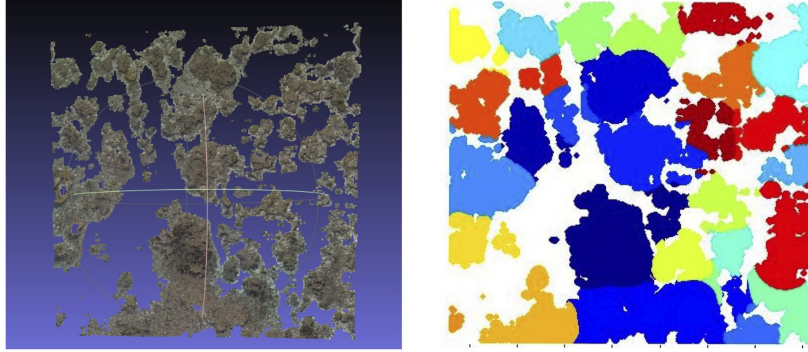


Figure 9: Clustering performed on the segmented output generated by approach 1 as shown in Fig. 8.

Figures 5 and 8 show the segmented point clouds corresponding to regions containing corals, as identified using the two different approaches outlined in the previous section. The next step in the process is to separate these point clouds into individual clusters that each represent one coral, to build a 3D model of each individual coral. As one can imagine, this is a slightly poorly defined task since it is hard to ascertain where one coral ends where another one begins. Sometimes, corals of one type can grow on top of dead corals of a different type as well.

We used the mean shift clustering [5] algorithm based on Euclidean distance of the points. Figs. 9 and 10 show the results of clustering algorithm. As one can see, while

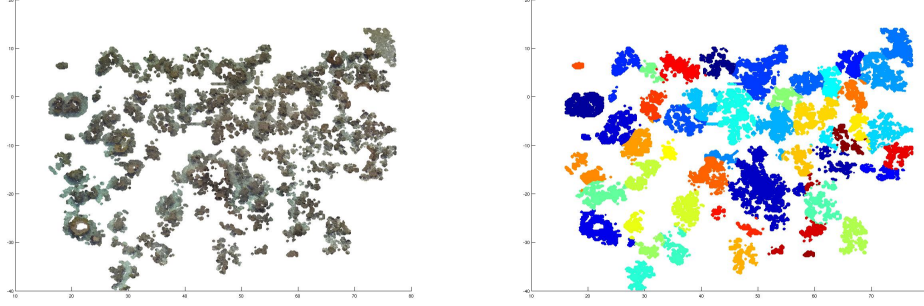


Figure 10: Clustering performed on the segmented output generated by approach 2 as shown in Fig. 5.

the output could have been better, the mean-shift algorithm does a fair job of dividing the corals into individual clusters. The width of the kernel used in the algorithm is an important parameter that needs to be tuned to achieve the best results. If the parameter is chosen to be too large, it ends up merging the corals into macro-clusters and if the parameter is too small, then it ends up generating a large number of small clusters, often breaking up a single coral into two separate clusters, especially when clustering the segmentation results from approach 2.

In order to verify that at the end of this process we have identified the clusters related to or containing corals, we showed each of the 3D point cloud clusters to the researchers working in this field. They were able to verify that the clusters actually contained the corals.

### 2.5. Classification

The biologists are interested in finding the total mass of different coral species in the reef in order to obtain a true picture of the health of the reef. Hence in order to classify the corals, an SVM based classifier was developed using the ‘radial basis function’ kernel. The two main types of corals present on the reef are

- Branching
- Porites



a.Branching Coral

b.Porites Coral



We used an SVM with the 'radial basis function' kernel to generate the classifier. Feature vectors were constructed for each of the clusters using

- height distribution
- curvature distribution
- color

Due to the sheer number of points in the point cloud and the limited computational resources at hand, we were able to generate only 128 clusters of corals. We hence had to keep the size of our feature vector small and to compensate for the small size of training data, k-fold cross validation was used to test the performance of the classifier.

### 3. Results

The results of segmentation and clustering are summarized in the form of precision, recall and F score values in Table 1. In order to obtain this table, a cluster largely containing the sea floor is defined to be a false positive and a cluster devoid of a coral is defined as false negative. Since these estimates are subjective, they should only be looked upon as approximate indicators of the success of our algorithm.

Table 1: Precision and Recall values for the two approaches

|   | Approach 1 (Chan Vese) | Approach 2 (Morphological Snakes) |
|---|------------------------|-----------------------------------|
| P | 0.86                   | 0.61                              |
| R | 0.93                   | 0.91                              |
| F | 0.89                   | 0.73                              |

For classification, we were able to achieve 75% accuracy using just the color information and using the 3D information in the form of height and curvature improved the accuracy to 76.6%. While the increment is not drastic, it shows that there is potential in improving the performance of traditional 2D-classification techniques using 3D information.

### 4. Concluding Remarks

We have developed a novel pipeline to go from 2D aerial images to segmented and classified clusters of 3D point clouds which allows the users to come up with estimates of volume and extract other useful pieces of information over large regions of interest. We used a commercial Structure From Motion toolbox to generate a 3D point cloud from the images. A three dimensional surface was fitted onto the point cloud which was then segmented into regions containing corals and regions devoid of them using two different techniques. We found that using the mean curvature of the 3D surface can add significant value to the segmentation output and there is hence significant potential in combining the 2D and 3D information based segmentation techniques to generate better segmentation

results. The coral containing segment was further subjected to clustering using the mean-shift algorithm. While the clustering output was not ideal, it did a fair job at giving us individual clusters of corals which were then classified using a combination of 2D and 3D features. While each of the steps in this pipeline can be tweaked and improved to get better results, our project provides a framework for extracting semantics and 3D information from 2D aerial images. We have shown that using 3D information has the potential to improve the performance of both segmentation and classification tasks.

This method can be easily extended to develop a multi-class classifier to distinguish between the corals and other objects such as rocks covered in algae, and in the presence of more training data, robust feature vectors incorporating more of the 3D information can be developed. In particular for images such as the images of corals that we have dealt with, standard 2D descriptors such as SIFT and SURF might not perform well due to the presence of repeating patterns and the absence of sharp features for those algorithms to latch on to. We can hence augment such descriptors with 3D information to obtain a more robust and comprehensive descriptor. This methodology can be applied to a wide variety of problems involving aerial imaging such as monitoring deforestation, disaster management etc.

## 5. References

- [1] T. Chan, L. Vese, Active contours without edges, *Image Processing, IEEE Transactions on* 10 (2) (2001) 266–277. doi:10.1109/83.902291.
- [2] M. Kazhdan, M. Bolitho, H. Hoppe, Poisson surface reconstruction, in: *Proceedings of the Fourth Eurographics Symposium on Geometry Processing, SGP '06*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2006, pp. 61–70.  
URL <http://dl.acm.org/citation.cfm?id=1281957.1281965>
- [3] CGAL, Computational Geometry Algorithms Library, <http://www.cgal.org>.
- [4] L. Alvarez, L. Baumela, P. Henriquez, P. Marquez-Neila, Morphological snakes, 2013 IEEE Conference on Computer Vision and Pattern Recognition 0 (2010) 2197–2202. doi:http://doi.ieeecomputersociety.org/10.1109/CVPR.2010.5539900.
- [5] D. Comaniciu, P. Meer, S. Member, Mean shift: A robust approach toward feature space analysis, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (2002) 603–619.